

# The Complexity of Online Manipulation of Sequential Elections\*

**Edith Hemaspaandra<sup>†</sup>**

Dept. of Computer Science  
Rochester Institute of Technology  
Rochester, NY 14623, USA

**Lane A. Hemaspaandra<sup>‡</sup>**

Dept. of Computer Science  
University of Rochester  
Rochester, NY 14627, USA

**Jörg Rothe<sup>§</sup>**

Inst. für Informatik  
Heinrich-Heine-Univ. Düsseldorf  
40225 Düsseldorf, Germany

## Abstract

Most work on manipulation assumes that all preferences are known to the manipulators. However, in many settings elections are open and sequential, and manipulators may know the already cast votes but may not know the future votes. We introduce a framework, in which manipulators can see the past votes but not the future ones, to model online coalitional manipulation of sequential elections, and we show that in this setting manipulation can be extremely complex even for election systems with simple winner problems. Yet we also show that for some of the most important election systems such manipulation is simple in certain settings. This suggests that when using sequential voting, one should pay great attention to the details of the setting in choosing one's voting rule.

Among the highlights of our classifications are: We show that, depending on the size of the manipulative coalition, the online manipulation problem can be complete for each level of the polynomial hierarchy or even for PSPACE. And we obtain the most dramatic contrast to date between the nonunique-winner and unique-winner models: Online weighted manipulation for plurality is in P in the nonunique-winner model, yet is coNP-hard (constructive case) and NP-hard (destructive case) in the unique-winner model.

## Introduction

Voting is a widely used method for preference aggregation and decision-making. In particular, *strategic voting* (or *manipulation*) has been studied intensely in social choice theory (starting with the celebrated work of Gibbard (1973) and Satterthwaite (1975)) and, in the rapidly emerging area of *computational social choice*, also with respect to its algorithmic properties and computational complexity (starting with the seminal work of Bartholdi, Tovey, and Trick (1989); see

\*Also appears as URCS-TR-2012-974.

<sup>†</sup>Supported in part by grant NSF-CCF-1101452 and a Friedrich Wilhelm Bessel Research Award. Work done in part while visiting Heinrich-Heine-Universität Düsseldorf.

<sup>‡</sup>Supported in part by grants NSF-CCF-{0426761,0915792, 1101479} and arc-dp110101792, and a Friedrich Wilhelm Bessel Research Award. Work done in part while visiting Heinrich-Heine-Universität Düsseldorf.

<sup>§</sup>Supported in part by DFG grant RO-1202/15-1, SFF grant “Cooperative Normsetting” of HHU Düsseldorf, and a DAAD grant for a PPP project in the PROCOPE program.

the recent surveys by Faliszewski and Procaccia (2010), Faliszewski, Hemaspaandra, and Hemaspaandra (2010), and Faliszewski et al. (2009)). This computational aspect is particularly important in light of the many applications of voting in computer science, ranging from meta-search heuristics for the internet (Dwork et al. 2001), to recommender systems (Ghosh et al. 1999) and multiagent systems in artificial intelligence (see the survey by Conitzer (2010)).

Most of the previous work on manipulation, however, is concerned with voting where the manipulators know the nonmanipulative votes. Far less attention has been paid (see the related work below) to manipulation in the midst of elections that are modeled as dynamic processes.

We introduce a novel framework for online manipulation, where voters vote in sequence and the current manipulator, who knows the previous votes and which voters are still to come but does not know their votes, must decide—right at that moment—what the “best” vote to cast is. So, while other approaches to sequential voting are game-theoretic, stochastic, or axiomatic in nature (again, see the related work), our approach to manipulation of sequential voting is shaped by the area of “online algorithms” (Borodin and El-Yaniv 1998), in the technical sense of a setting in which one (for us, each manipulative voter) is being asked to make a manipulation decision just on the basis of the information one has in one’s hands at the moment even though additional information/system evolution may well be happening down the line. In this area, there are different frameworks for evaluation. But the most attractive one, which pervades the area as a general theme, is the idea that one may want to “maxi-min” things—one may want to take the action that maximizes the goodness of the set of outcomes that one can expect regardless of what happens down the line from one time-wise. For example, if the current manipulator’s preferences are Alice > Ted > Carol > Bob and if she can cast a (perhaps insincere) vote that ensures that Alice or Ted will be a winner no matter what later voters do, and there is no vote she can cast that ensures that Alice will always be a winner, this maxi-min approach would say that that vote is a “best” vote to cast.

It will perhaps be a bit surprising to those familiar with online algorithms and competitive analysis that in our model of online manipulation we will not use a (competitive) ratio. The reason is that voting commonly uses an *ordinal* pref-

erence model, in which preferences are total orders of the candidates. It would be a severely improper step to jump from that to assumptions about intensity of preferences and utility, e.g., to assuming that everyone likes her  $n$ th-to-least favorite candidate exactly  $n$  times more than she likes her least favorite candidate.

**Related Work.** Xia and Conitzer (2010a) (see also the related paper by Desmedt and Elkind (2010)) define and study the Stackelberg voting game (also quite naturally called, in an earlier paper that mostly looked at two candidates, the roll-call voting game (Sloth 1993)). This basically is an election in which the voters vote in order, *and the preferences are common knowledge—everyone knows everyone else's preferences, everyone knows that everyone knows everyone else's preferences, and so on out to infinity*. Their analysis of this game is fundamentally game-theoretic; with such complete knowledge in a sequential setting, there is precisely one (subgame perfect Nash) equilibrium, which can be computed from the back end forward. Under their work's setting and assumptions, for bounded numbers of manipulators manipulation is in P, but we will show that in our model even with bounded numbers of manipulators manipulation sometimes (unless P = NP) falls beyond P.

The interesting “dynamic voting” work of Tennenholz (2004) investigates sequential voting, but focuses on axioms and voting rules rather than on coalitions and manipulation. Much heavily Markovian work studies sequential decision-making and/or dynamically varying preferences (see (Parkes and Procaccia 2011) and the references therein); our work in contrast is nonprobabilistic and focused on the complexity of coalitional manipulation. Also somewhat related to, but quite different from, our work is the work on possible and necessary winners. The seminal paper on that is due to Konczak and Lang (2005), and more recent work includes (Xia and Conitzer 2008; Betzler, Hemmann, and Niedermeier 2009; Bachrach, Betzler, and Faliszewski 2010; Betzler 2010; Betzler and Dorn 2010; Maudet et al. 2010; Baumeister and Rothe 2012); the biggest difference is that those are, loosely, one-quantifier settings, but the more dynamic setting of online manipulation involves numbers of quantifiers that can grow with the input size. Another related research line studies multi-issue elections (Xia and Conitzer 2010b; Xia, Conitzer, and Lang 2010; Xia, Conitzer, and Lang 2011; Xia, Lang, and Conitzer 2011); although there the separate issues may run in sequence, each issue typically is voted on simultaneously and with preferences being common knowledge.

## Preliminaries

**Elections.** A (*standard, i.e., simultaneous*) election  $(C, V)$  is specified by a set  $C$  of candidates and a list  $V$ , where we assume that each element in  $V$  is a pair  $(v, p)$  such that  $v$  is a voter name and  $p$  is  $v$ 's vote. How the votes in  $V$  are represented depends on the election system used—we assume, as is required by most systems, votes to be total preference orders over  $C$ . If, say,  $C = \{a, b, c\}$ , a vote of the form  $c > a > b$

means that this voter (strictly) prefers  $c$  to  $a$  and  $a$  to  $b$ .

We introduce election snapshots to capture sequential election scenarios as follows. Let  $C$  be a set of candidates and let  $u$  be (the name of) a voter. An *election snapshot for  $C$  and  $u$*  is specified by a triple  $V = (V_{<u}, u, V_{u<})$  consisting of all voters in the order they will vote, along with, for each voter before  $u$  (i.e., those in  $V_{<u}$ ), the vote she cast, and for each voter after  $u$  (i.e., those in  $V_{u<}$ ), a bit specifying if she is part of the manipulative coalition (to which  $u$  always belongs). That is,  $V_{<u} = ((v_1, p_1), (v_2, p_2), \dots, (v_{i-1}, p_{i-1}))$ , where the voters named  $v_1, v_2, \dots, v_{i-1}$  (including perhaps manipulators and nonmanipulators) have already cast their votes (preference order  $p_j$  being cast by  $v_j$ ), and  $V_{u<} = ((v_{i+1}, x_{i+1}), (v_{i+2}, x_{i+2}), \dots, (v_n, x_n))$  lists the names of the voters still to cast their votes, in that order, and where  $x_j = 1$  if  $v_j$  belongs to the manipulative coalition and  $x_j = 0$  otherwise.

**Scoring Rules.** A *scoring rule* for  $m$  candidates is given by a scoring vector  $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$  of nonnegative integers such that  $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$ . For an election  $(C, V)$ , each candidate  $c \in C$  scores  $\alpha_i$  points for each vote that ranks  $c$  in the  $i$ th position. Let  $score(c)$  be the total score of  $c \in C$ . All candidates scoring the most points are winners of  $(C, V)$ . Some of the most popular voting systems are *k*-*approval* (esp. *plurality*, aka 1-*approval*) and *k*-*veto* (esp. *veto*, aka 1-*veto*). Their  $m$ -candidate,  $m \geq k$ , versions are defined by the scoring vectors  $(\underbrace{1, \dots, 1}_{k}, \underbrace{0, \dots, 0}_{m-k})$  and  $(\underbrace{1, \dots, 1}_{m-k}, \underbrace{0, \dots, 0}_k)$ .

When  $m$  is not fixed, we omit the phrase “ $m$ -candidate.”

**Manipulation.** The (*standard*) *weighted coalitional manipulation problem* (Conitzer, Sandholm, and Lang 2007), abbreviated by  $\mathcal{E}$ -WCM, for any election system  $\mathcal{E}$  is as follows: Given a candidate set  $C$ , a list  $S$  of nonmanipulative voters each having a nonnegative integer weight, a list  $T$  of the nonnegative integer weights of the manipulative voters (whose preferences over  $C$  are unspecified), with  $S \cap T = \emptyset$ , and a distinguished candidate  $c \in C$ , can the manipulative votes  $T$  be set such that  $c$  is a (or the)  $\mathcal{E}$  winner of  $(C, S \cup T)$ ? Asking whether  $c$  can be made “a winner” is called the nonunique-winner model and is the model of all notions in this paper unless mentioned otherwise. If one asks whether  $c$  can be made a “one and only winner,” that is called the unique-winner model. We also use the *unweighted* variant, where each vote has unit weight, and write  $\mathcal{E}$ -UCM as a shorthand. Note that  $\mathcal{E}$ -UCM with a *single* manipulator (i.e.,  $\|T\| = 1$  in the problem instance) is the manipulation problem originally studied in (Bartholdi, Tovey, and Trick 1989; Bartholdi and Orlin 1991). Conitzer, Sandholm, and Lang (2007) also introduced the *destructive* variants of these manipulation problems, where the goal is not to make  $c$  win but to ensure that  $c$  is not a winner, and we denote the corresponding problems by  $\mathcal{E}$ -DWCM and  $\mathcal{E}$ -DUCM. Finally, we write  $\mathcal{E}$ -WC $\neq 0$ M,  $\mathcal{E}$ -UC $\neq 0$ M,  $\mathcal{E}$ -DWC $\neq 0$ M, and  $\mathcal{E}$ -DUC $\neq 0$ M to indicate that the problem instances are required to have a nonempty coalition of manipulators.

**Complexity-Theoretic Background.** We assume the reader is familiar with basic complexity-theoretic notions

such as the complexity classes P and NP, the class FP of polynomial-time computable functions, polynomial-time many-one reducibility ( $\leq_m^p$ ), and hardness and completeness with respect to  $\leq_m^p$  for a complexity class.

Meyer and Stockmeyer (1972) and Stockmeyer (1976) introduced and studied the polynomial hierarchy,  $\text{PH} = \bigcup_{k \geq 0} \Sigma_k^p$ , whose levels are inductively defined by  $\Sigma_0^p = \text{P}$  and  $\Sigma_{k+1}^p = \text{NP}^{\Sigma_k^p}$ , and their co-classes,  $\Pi_k^p = \text{co}\Sigma_k^p$  for  $k \geq 0$ . They also characterized these levels by polynomially length-bounded alternating existential and universal quantifiers.

$\text{P}^{\text{NP}}$  is the class of problems solvable in deterministic polynomial time with access to an NP oracle, and  $\text{P}^{\text{NP}[1]}$  is the restriction of  $\text{P}^{\text{NP}}$  where only one oracle query is allowed. Note that  $\text{P} \subseteq \text{NP} \cup \text{coNP} \subseteq \text{P}^{\text{NP}[1]} \subseteq \text{P}^{\text{NP}} \subseteq \Sigma_2^p \cap \Pi_2^p \subseteq \Sigma_2^p \cup \Pi_2^p \subseteq \text{PH} \subseteq \text{PSPACE}$ , where PSPACE is the class of problems solvable in polynomial space. The *quantified boolean formula problem*, QBF, is a standard PSPACE-complete problem. Define  $\text{QBF}_k$  ( $\widetilde{\text{QBF}}_k$ ) to be the restriction of QBF with at most  $k$  quantifiers that start with  $\exists$  ( $\forall$ ) and then alternate between  $\exists$  and  $\forall$ , and we assume that each  $\exists$  and  $\forall$  quantifies over a set of boolean variables. For each  $k \geq 1$ ,  $\text{QBF}_k$  is  $\Sigma_k^p$ -complete and  $\widetilde{\text{QBF}}_k$  is  $\Pi_k^p$ -complete.

## Our Model of Online Manipulation

The core of our model of online manipulation in sequential voting is what we call the *magnifying-glass moment*, namely, the moment at which a manipulator  $u$  is the one who is going to vote, is aware of what has happened so far in the election (and which voters are still to come, but in general not knowing what they want, except in the case of voters, if any, who are coalitionally linked to  $u$ ). In this moment,  $u$  seeks to “figure out” what the “best” vote to cast is. We will call the information available in such a moment an *online manipulation setting* (OMS, for short) and define it formally as a tuple  $(C, u, V, \sigma, d)$ , where  $C$  is a set of candidates;  $u$  is a distinguished voter;  $V = (V_{<u}, u, V_{u<})$  is an election snapshot for  $C$  and  $u$ ;  $\sigma$  is the preference order of the manipulative coalition to which  $u$  belongs; and  $d \in C$  is a distinguished candidate. Given an election system  $\mathcal{E}$ , define the *online unweighted coalitional manipulation problem*, abbreviated by online- $\mathcal{E}$ -UCM, as follows: Given an OMS  $(C, u, V, \sigma, d)$  as described above, does there exist some vote that  $u$  can cast (assuming support from the manipulators coming after  $u$ ) such that no matter what votes are cast by the nonmanipulators coming after  $u$ , there exists some  $j \in C$  such that  $j \geq_\sigma d$  and  $j$  is an  $\mathcal{E}$  winner of the election?

By “support from the manipulators coming after  $u$ ” we mean that  $u$ ’s coalition partners coming after  $u$ , when they get to vote, will use their then-in-hand knowledge of all votes up to then to help  $u$  reach her goal: By a joint effort  $u$ ’s coalition can ensure that the  $\mathcal{E}$  winner set will always include a candidate liked by the coalition as much as or more than  $d$ , even when the nonmanipulators take their strongest action so as to prevent this. Note that this candidate,  $j$  in the problem description, may be different based on the nonmanipulators’ actions. (Nonsequential manipulation problems usually focus on whether a single candidate can be made to

win, but in our setting, this “that person or better” focus is more natural.)

For the case of weighted manipulation, each voter also comes with a nonnegative integer weight. We denote this problem by online- $\mathcal{E}$ -WCM.

We write online- $\mathcal{E}$ -UCM[ $k$ ] in the unweighted case and online- $\mathcal{E}$ -WCM[ $k$ ] in the weighted case to denote the problem when the number of manipulators from  $u$  onward is restricted to be at most  $k$ .

Our corresponding destructive problems are denoted by online- $\mathcal{E}$ -DUCM, online- $\mathcal{E}$ -DWCM, online- $\mathcal{E}$ -DUCM[ $k$ ], and online- $\mathcal{E}$ -DWCM[ $k$ ]. In online- $\mathcal{E}$ -DUCM we ask whether the given current manipulator  $u$  (assuming support from the manipulators after her) can cast a vote such that no matter what votes are cast by the nonmanipulators after  $u$ , no  $j \in C$  with  $d \geq_\sigma j$  is an  $\mathcal{E}$  winner of the election, i.e.,  $u$ ’s coalition can ensure that the  $\mathcal{E}$  winner set never includes  $d$  or any even more hated candidate. The other three problems are defined analogously.

Note that online- $\mathcal{E}$ -UCM generalizes the original unweighted manipulation problem with a single manipulator as introduced by Bartholdi, Tovey, and Trick (1989). Indeed, their manipulation problem in effect is the special case of online- $\mathcal{E}$ -UCM when restricted to instances where there is just one manipulator, she is the last voter to cast a vote, and  $d$  is the coalition’s most preferred candidate. Similarly, online- $\mathcal{E}$ -WCM generalizes the (standard) coalitional weighted manipulation problem (for nonempty coalitions of manipulators). Indeed, that traditional manipulation problem is the special case of online- $\mathcal{E}$ -WCM, restricted to instances where only manipulators come after  $u$  and  $d$  is the coalition’s most preferred candidate. If we take an analogous approach except with  $d$  restricted now to being the most hated candidate of the coalition, we generalize the corresponding notions for the destructive cases. We summarize these observations as follows.

**Proposition 1** *For each election system  $\mathcal{E}$ , it holds that (a)  $\mathcal{E}\text{-UC}_{\neq\emptyset}M \leq_m^p$  online- $\mathcal{E}$ -UCM, (b)  $\mathcal{E}\text{-WC}_{\neq\emptyset}M \leq_m^p$  online- $\mathcal{E}$ -WCM, (c)  $\mathcal{E}\text{-DUC}_{\neq\emptyset}M \leq_m^p$  online- $\mathcal{E}$ -DUCM, and (d)  $\mathcal{E}\text{-DW}_{\neq\emptyset}M \leq_m^p$  online- $\mathcal{E}$ -DWCM.*

**Corollary 2** *For each election system  $\mathcal{E}$  such that the winner problem is solvable in polynomial time, it holds that (a)  $\mathcal{E}\text{-UCM} \leq_m^p$  online- $\mathcal{E}$ -UCM, (b)  $\mathcal{E}\text{-WCM} \leq_m^p$  online- $\mathcal{E}$ -WCM, (c)  $\mathcal{E}\text{-DUCM} \leq_m^p$  online- $\mathcal{E}$ -DUCM, and (d)  $\mathcal{E}\text{-DWCM} \leq_m^p$  online- $\mathcal{E}$ -DWCM.*

We said above that, by default, we will use the *nonunique-winner model* and all the above problems are defined in this model. However, we will also have some results in the *unique-winner model*, which will, here, sharply contrast with the corresponding results in the nonunique-winner model. To indicate that a problem, such as online- $\mathcal{E}$ -UCM, is in the unique-winner model, we write online- $\mathcal{E}$ -UCM<sub>UW</sub> and ask whether the current manipulator  $u$  (assuming support from the manipulators coming after her) can ensure that there exists some  $j \in C$  such that  $j \geq_\sigma d$  and  $j$  is the *unique  $\mathcal{E}$  winner* of the election.

## General Results

**Theorem 3 1.** For each election system  $\mathcal{E}$  whose winner problem can be solved in polynomial time, online- $\mathcal{E}$ -WCM is in PSPACE (and so is online- $\mathcal{E}$ -UCM).

2. There exists an election system  $\mathcal{E}$  with a polynomial-time winner problem such that online- $\mathcal{E}$ -UCM is PSPACE-complete (and so is online- $\mathcal{E}$ -WCM).

PROOF. We omit the easy proof of the first statement. We construct an election system  $\mathcal{E}$  establishing the second statement. Given an input  $(C, u, V, \sigma, d)$ ,  $\mathcal{E}$  will look at the lexicographically least candidate name in  $C$ . Let  $c$  represent that name string in some fixed, natural encoding.  $\mathcal{E}$  will check if  $c$  represents a *tiered* boolean formula, by which we mean one whose variable names are all of the form  $x_{i,j}$  (which really means a direct encoding of a string, such as “ $x_{4,9}$ ”); the  $i, j$  fields must all be positive integers. If  $c$  does not represent such a tiered formula, everyone loses on that input. Otherwise (i.e., if  $c$  represents a tiered formula), let  $width$  be the maximum  $j$  occurring as the second subscript in any variable name  $(x_{i,j})$  in  $c$ , and let  $blocks$  be the maximum  $i$  occurring as the first subscript in any variable name in  $c$ . If there are fewer than  $blocks$  voters in  $V$ , everyone loses. Otherwise, if there are fewer than  $1 + 2 \cdot width$  candidates in  $C$ , everyone loses (this is so that each vote will involve enough candidates that it can be used to set all the variables in one block). Otherwise, if there exists some  $i, 1 \leq i \leq blocks$ , such that for no  $j$  does the variable  $x_{i,j}$  occur in  $c$ , then everyone loses. Otherwise, order the voters from the lexicographically least to the lexicographically greatest voter name. Now, the first voter in this order will assign truth values to all variables  $x_{1,*}$ , the second voter in this order will assign truth values to all variables  $x_{2,*}$ , and so on up to the  $blockst$  voter, who will assign truth values to all variables  $x_{blocks,*}$ .

How do we get those assignments from these votes? Consider a vote whose total order over  $C$  is  $\sigma'$  (and recall that  $\|C\| \geq 1 + 2 \cdot width$ ). Remove  $c$  from  $\sigma'$ , yielding  $\sigma''$ . Let  $c_1 <_{\sigma''} c_2 <_{\sigma''} \dots <_{\sigma''} c_{2 \cdot width}$  be the  $2 \cdot width$  least preferred candidates in  $\sigma''$ . We build a vector in  $\{0, 1\}^{width}$  as follows: The  $\ell$ th bit of the vector is 0 if the string that names  $c_{1+2(\ell-1)}$  is lexicographically less than the string that names  $c_{2\ell}$ , and this bit is 1 otherwise.

Let  $b_i$  denote the vector thus built from the  $i$ th vote (in the above ordering),  $1 \leq i \leq blocks$ . Now, for each variable  $x_{i,j}$  occurring in  $c$ , assign to it the value of the  $j$ th bit of  $b_i$ , where 0 represents *false* and 1 represents *true*. We have now assigned all variables of  $c$ , so  $c$  evaluates to either *true* or *false*. If  $c$  evaluates to *true*, everyone wins, otherwise everyone loses. This completes the specification of the election system  $\mathcal{E}$ .  $\mathcal{E}$  has a polynomial-time winner problem, as any boolean formula, given an assignment to all its variables, can easily be evaluated in polynomial time.

To show PSPACE-hardness, we  $\leq_m^p$ -reduce the PSPACE-complete problem QBF to online- $\mathcal{E}$ -UCM. Let  $y$  be an instance of QBF. We transform  $y$  into an instance of the form

$$(\exists x_{1,1}, \dots, x_{1,k_1}) (\forall x_{2,1}, \dots, x_{2,k_2}) \cdots (Q_\ell x_{\ell,1}, \dots, x_{\ell,k_\ell})$$

$$[\Phi(x_{1,1}, \dots, x_{1,k_1}, x_{2,1}, \dots, x_{2,k_2}, \dots, x_{\ell,1}, \dots, x_{\ell,k_\ell})]$$

in polynomial time, where  $Q_\ell = \exists$  if  $\ell$  is odd and  $Q_\ell = \forall$  if  $\ell$  is even, the  $x_{i,j}$  are boolean variables,  $\Phi$  is a boolean

formula, and for each  $i, 1 \leq i \leq \ell$ ,  $\Phi$  contains at least one variable of the form  $x_{i,*}$ . This quantified boolean formula is  $\leq_m^p$ -reduced to an instance  $(C, u, V, \sigma, c)$  of online- $\mathcal{E}$ -UCM as follows:

1.  $C$  contains a candidate whose name,  $c$ , encodes  $\Phi$ , and in addition  $C$  contains  $2 \cdot \max(k_1, \dots, k_\ell)$  other candidates, all with names lexicographically greater than  $c$ —for specificity, let us say their names are the  $2 \cdot \max(k_1, \dots, k_\ell)$  strings that immediately follow  $c$  in lexicographic order.
2.  $V$  contains  $\ell$  voters,  $1, 2, \dots, \ell$ , who vote in that order, where  $u = 1$  is the distinguished voter and all odd voters belong to  $u$ 's manipulative coalition and all even voters do not. The voter names will be lexicographically ordered by their number, 1 is least and  $\ell$  is greatest.
3. The manipulators' preference order  $\sigma$  is to like candidates in the opposite of their lexicographic order. In particular,  $c$  is the coalition's most preferred candidate.

This is a polynomial-time reduction. It follows immediately from this construction and the definition of  $\mathcal{E}$  that  $y$  is in QBF if and only if  $(C, u, V, \sigma, c)$  is in online- $\mathcal{E}$ -UCM.  $\square$

The following theorem shows that for bounded numbers of manipulators the complexity crawls up the polynomial hierarchy. The theorem's (omitted) proof is based on the proof given above, except we need to use the alternating quantifier characterization due to Meyer and Stockmeyer (1972) and Stockmeyer (1976) for the upper bound and to reduce from the  $\Sigma_{2k}^p$ -complete problem QBF<sub>2k</sub> rather than from QBF for the lower bound.

**Theorem 4** Fix any  $k \geq 1$ . 1. For each election system  $\mathcal{E}$  whose winner problem can be solved in polynomial time, online- $\mathcal{E}$ -WCM[k] is in  $\Sigma_{2k}^p$  (and so is online- $\mathcal{E}$ -UCM[k]).

2. There exists an election system  $\mathcal{E}$  with a polynomial-time winner problem such that online- $\mathcal{E}$ -UCM[k] is  $\Sigma_{2k}^p$ -complete (and so is online- $\mathcal{E}$ -WCM[k]).

Note that the (constructive) online manipulation problems considered in Theorems 3 and 4 are about ensuring that the winner set always contains some candidate in the  $\sigma$  segment stretching from  $d$  up to the top-choice. Now consider “pinpoint” variants of these problems, where we ask whether the distinguished candidate  $d$  herself can be guaranteed to be a winner (for nonsequential manipulation, that version indeed is the one commonly studied). Denote the *pinpoint* variant of, e.g., online- $\mathcal{E}$ -UCM[k] by pinpoint-online- $\mathcal{E}$ -UCM[k]. Since our hardness proofs in Theorems 3 and 4 make all or no one a winner (and as the upper bounds in these theorems also can be seen to hold for the pinpoint variants), they establish the corresponding completeness results also for the pinpoint cases. We thus have completeness results for PSPACE and  $\Sigma_{2k}^p$  for each  $k \geq 1$ . What about the classes  $\Sigma_{2k+1}^p$  and  $\Pi_k^p$ , for each  $k \geq 1$ ? We can get completeness results for all these classes by defining appropriate variants of online manipulation problems. Let OMP be any of the online manipulation problems considered earlier, including the pinpoint variants mentioned above. Define freeform-OMP to be just as OMP, except we no longer require the distinguished voter  $u$  to be part of the manipulative coalition— $u$  can be in or can be

out, and the input must specify, for  $u$  and all voters in  $V_{u <}$ , which ones are the members of the coalition. The question of freeform-OMP is whether it is true that for all actions of the nonmanipulators at or after  $u$  (for specificity as to this problem:  $u$  if a nonmanipulator will in the input come with a preference order) there will be actions (each taken with full information on cast-before-them votes) of the manipulative coalition members such that their goal of making some candidate  $j$  with  $j \geq_\sigma d$  (or exactly  $d$ , in the pinpoint versions) a winner is achieved. Then, whenever Theorem 4 establishes a  $\Sigma_{2k}^p$  or  $\Sigma_{2k}^p$ -completeness result for OMP, we obtain a  $\Pi_{2k+1}^p$  or  $\Pi_{2k+1}^p$ -completeness result for freeform-OMP. Similarly, the PSPACE and PSPACE-completeness results for OMP we established in Theorem 3 also can be shown true for freeform-OMP. On the other hand, if we define a variant of OMP by requiring the final voter to always be a manipulator, the PSPACE and PSPACE-completeness results for OMP from Theorem 3 remain true for this variant; the  $\Sigma_{2k}^p$  and  $\Sigma_{2k}^p$ -completeness results for OMP from Theorem 4 change to  $\Sigma_{2k-1}^p$  and  $\Sigma_{2k-1}^p$ -completeness results for this variant; and the above  $\Pi_{2k+1}^p$  and  $\Pi_{2k+1}^p$ -completeness results for freeform-OMP change to  $\Pi_{2k}^p$  and  $\Pi_{2k}^p$ -completeness results for this variant.

Finally, as an open direction (and related conjecture), we define for each of the previously considered variants of online manipulation problems a *full profile* version. For example, fullprofile-online- $\mathcal{E}$ -UCM[ $k$ ] is the function problem for a given election system  $\mathcal{E}$  that, given an OMS without any distinguished candidate,  $(C, u, V, \sigma)$ , returns a length  $\|C\|$  bit-vector that for each candidate  $d \in C$  says if the answer to “ $(C, u, V, \sigma, d) \in \text{online-}\mathcal{E}\text{-UCM}[k]$ ” is “yes” (1) or “no” (0). The function problem fullprofile-pinpoint-online- $\mathcal{E}$ -UCM[ $k$ ] is defined analogously, except regarding pinpoint-online- $\mathcal{E}$ -UCM[ $k$ ].

**Theorem 5 1.** fullprofile-online- $\mathcal{E}$ -UCM[ $k$ ] is in  $\text{FP}^{\Sigma_{2k}^p[\mathcal{O}(\log n)]}$ , the class of functions computable in polynomial time given Turing access to a  $\Sigma_{2k}^p$  oracle with  $\mathcal{O}(\log n)$  queries allowed on inputs of size  $n$ .

2. fullprofile-pinpoint-online- $\mathcal{E}$ -UCM[ $k$ ] is in  $\text{FP}_{tt}^{\Sigma_{2k}^p}$ , the class of functions computable in polynomial time given truth-table access to a  $\Sigma_{2k}^p$  oracle.

We conjecture that both problems are complete for the corresponding class under metric reductions (Krentel 1988).

## Results for Specific Natural Voting Systems

The results of the previous section show that, simply put, even for election systems with polynomial-time winner problems, online manipulation can be tremendously difficult. But what about *natural* election systems? We will now take a closer look at important natural systems. We will show that online manipulation can be easy for them, depending on which particular problem is considered, and we will also see that the constructive and destructive cases can differ sharply from each other and that it really matters whether we are in the nonunique-winner model or the unique-winner model. The proof of Theorem 6 is omitted due to space limitations.

**Theorem 6 1.** fullprofile-online-plurality-WCM (and thus also fullprofile-online-plurality-UCM) is in FP.

2. fullprofile-online-plurality-DWCM (and thus also fullprofile-online-plurality-DUCM) is in FP.

**Corollary 7** The problems online-plurality-WCM (and thus also online-plurality-UCM) and online-plurality-DWCM (and thus also online-plurality-DUCM) are each in P.

Theorem 6 and Corollary 7 refer to problems in the nonunique-winner model. By contrast, we now show that online manipulation for weighted plurality voting in the *unique-winner* model is coNP-hard in the *constructive* case and is NP-hard in the *destructive* case. This is perhaps the most dramatic, broad contrast yet between the nonunique-winner model and the unique-winner model, and is the first such contrast involving plurality. The key other NP-hardness versus P result for the nonunique-winner model versus the unique-winner model is due to Faliszewski, Hemaspaandra, and Schnoor (2008), but holds only for (standard) weighted manipulation for Copeland elections with exactly three candidates; for fewer than three both cases there are in P and for more than three both are NP-complete. In contrast, the P results of Corollary 7 hold for all numbers of candidates, and the NP-hardness and coNP-hardness results of Theorem 8 hold whenever there are at least two candidates.

**Theorem 8 1.** The problem online-plurality-WCM<sub>UW</sub> is coNP-hard, even when restricted to only two candidates (and this also holds when restricted to three, four, ... candidates).

2. The problem online-plurality-DWCM<sub>UW</sub> is NP-hard, even when restricted to only two candidates (and this also holds when restricted to three, four, ... candidates).

Due to space limitations, we do not present a full proof of Theorem 8 but only sketch the proof of NP-hardness for the second part, by a reduction from the NP-complete problem Partition: Given a nonempty sequence  $(w_1, w_2, \dots, w_z)$  of positive integers such that  $\sum_{i=1}^z w_i = 2W$ , does there exist a set  $I \subseteq \{1, 2, \dots, z\}$  such that  $\sum_{i \in I} w_i = W$ ? Let  $m \geq 2$ . Given an instance  $(w_1, w_2, \dots, w_z)$  of Partition, construct an instance  $((\{c_1, \dots, c_m\}, u_1, V, c_1 > c_2 > \dots > c_m, c_1))$  of online-plurality-DWCM<sub>UW</sub> such that  $V$  contains  $m+z-2$  voters  $v_1, \dots, v_{m-2}, u_1, \dots, u_z$  who vote in that order, for  $1 \leq i \leq m-2$ ,  $v_i$  votes for  $c_i$  and has weight  $(m-1)W - i$ , and for  $1 \leq i \leq z$ ,  $u_i$  is a manipulator of weight  $(m-1)w_i$ . If  $(w_1, w_2, \dots, w_z)$  is a yes-instance of Partition, the manipulators can give  $(m-1)W$  points to both  $c_{m-1}$  and  $c_m$ , and zero points to the other candidates. So  $c_{m-1}$  and  $c_m$  are tied for the most points and there is no unique winner. On the other hand, the only way to avoid having a unique winner in our online-plurality-DWCM<sub>UW</sub> instance is if there is a tie for the most points. The only candidates that can tie are  $c_{m-1}$  and  $c_m$ , since all other pairs of candidates have different scores modulo  $m-1$ . It is easy to see that  $c_{m-1}$  and  $c_m$  tie for the most points only if they both get exactly  $(m-1)W$  points. It follows that  $(w_1, w_2, \dots, w_z)$  is a yes-instance of Partition.

**Theorem 9** For each scoring rule  $\alpha = (\alpha_1, \dots, \alpha_m)$ , online- $\alpha$ -WCM is in P if  $\alpha_2 = \alpha_m$  and is NP-hard otherwise.

This follows from Corollaries 2 and 7 and the main theorem from Hemaspandra and Hemaspandra (2007).

**Theorem 10** For each  $k$ , online- $k$ -approval-UCM and online- $k$ -veto-UCM are in P.

**PROOF SKETCH.** Consider 1-veto. Given an online-1-veto-UCM instance  $(C, u, V, \sigma, d)$ , the best strategy for the manipulators from  $u$  onward (let  $n_1$  denote how many of these there are) is to minimize  $\max_{c < \sigma d} \text{score}(c)$ . Let  $n_0$  denote how many nonmanipulators come after  $u$ . We claim that  $(C, u, V, \sigma, d)$  is a yes-instance if and only if  $d$  is ranked last in  $\sigma$  or there exists a  $t$  such that (1)  $\sum_{c < \sigma d} (\text{maxscore}(c) \ominus t) \leq n_1$  (so those manipulators can ensure that all candidates ranked  $<_\sigma d$  score at most  $t$  points), where “ $\ominus$ ” denotes proper subtraction ( $x \ominus y = \max(x - y, 0)$ ) and  $\text{maxscore}(c)$  is  $c$ 's score when none of the voters from  $u$  onward veto  $c$ , and (2)  $\sum_{c \geq \sigma d} (\text{maxscore}(c) \ominus (t - 1)) > n_0$  (so those nonmanipulators cannot prevent that some candidate ranked  $\geq_\sigma d$  scores at least  $t$  points).

For 1-veto, under the above approach, in each situation where the remaining manipulators can force success against all actions of the remaining nonmanipulators,  $u$  (right then as she moves) can set her *and all future manipulators' actions* so as to force success regardless of the actions of the remaining nonmanipulators. For  $k$ -approval and  $k$ -veto,  $k \geq 2$ , that approach provably cannot work; rather, we sometimes need later manipulators' actions to be shaped by intervening nonmanipulators' actions. Still, the following P-time algorithm, which works for all  $k$ , tells if success can be forced. As a thought experiment, for each voter  $v$  from  $u$  onwards in sequence do this: Order the candidates in  $\{c \mid c \geq_\sigma d\}$  from most to least current approvals, and postpend the remaining candidates ordered from least to most current approvals. Let  $\ell$  be  $k$  for  $k$ -approval and  $\|C\| - k$  for  $k$ -veto. Cast the voter's  $\ell$  approvals for the first  $\ell$  candidates in this order if  $v$  is a manipulator and otherwise for the last  $\ell$  candidates in this order. Success can be forced against perfect play if and only if this P-time process leads to success.  $\square$

We now turn to 3-candidate-online-veto-WCM. Let  $C = \{a, b, c\}$  and suppose  $u$ 's manipulative coalition has the preference order  $a >_\sigma b >_\sigma c$ . Let  $d$  denote the distinguished candidate. The  $d = c$  inputs have a trivial P algorithm. Restricted to the  $d = a$  inputs, the problem is NP-hard, which follows from NP-hardness of 3-candidate-veto-WCM due to Conitzer, Sandholm, and Lang (2007) (who note their result is valid in the unique-winner and nonunique-winner models) and Corollary 2. This case is also in NP, by the following NP algorithm: Given an instance  $(C, u, V, \sigma, d)$  of 3-candidate-online-veto-WCM satisfying  $d >_\sigma b >_\sigma c$ , nondeterministically guess a partition  $(A, B)$  of the manipulators from  $u$  onward; all voters in  $A$  veto  $b$  and all voters in  $B$  veto  $c$ ; the nonmanipulators after  $u$  veto  $a = d$ ; on any such path, accept if and only if  $a$  is a winner. Restricted to  $d = b$ , 3-candidate-online-veto-WCM is coNP-hard, which

follows by a reduction from Partition to the complement of 3-candidate-online-veto-WCM. Observe that this case is also in coNP, by an NP algorithm for the complement similar to the above. This (since 3-candidate-online-veto-WCM can in light of the above be written as the union of an NP-complete and a coNP-complete set that are P-separable<sup>1</sup>) proves:

**Theorem 11** 3-candidate-online-veto-WCM is  $P^{NP[1]}$ -complete.

Dropping the restriction to three candidates, we obtain the following result, the proof of which is omitted, which places this problem far below the general PSPACE bound from earlier in this paper.

**Theorem 12** online-veto-WCM is in  $P^{NP}$ .

Immediately from Theorems 10 and 12, the full profile variants of online- $k$ -veto-UCM and online- $k$ -approval-UCM are in FP and fullprofile-online-veto-WCM is in  $FP^{NP}$ .

## Uncertainty About the Order of Future Voters

So far, we have been dealing with cases where the order of future voters was fixed and known. But what happens if the order of future voters itself is unknown? Even here, we can make claims. To model this most naturally, our “magnifying-glass moment” will focus not on one manipulator  $u$ , but will focus at a moment in time when some voters are still to come (as before, we know who they are and which are manipulators; as before, we have a preference order  $\sigma$ , and know what votes have been cast so far, and have a distinguished candidate  $d$ ). And the question our problem is asking is: Is it the case that our manipulative coalition can ensure that the winner set will always include  $d$  or someone liked more than  $d$  with respect to  $\sigma$  (i.e., the winner set will have nonempty intersection with  $\{c \in C \mid c \geq_\sigma d\}$ ), *regardless of what order the remaining voters vote in*. We will call this problem the *schedule-robust online manipulation problem*, and will denote it by SR-online- $\mathcal{E}$ -UCM. (We will add a “[1,1]” suffix for the restriction of this problem to instances when at most one manipulator and at most one nonmanipulator have not yet voted.) One might think that this problem captures both a  $\Sigma_2^P$  and a  $\Pi_2^P$  issue, and so would be hard for both classes. However, the requirement of schedule robustness tames the problem (basically what underpins that is simply that exists-forall-predicate implies forall-exists-predicate), bringing it into  $\Sigma_2^P$ . Further, we can prove, by explicit construction of such a system, that for some simple election systems this problem is complete for  $\Sigma_2^P$ .

**Theorem 13 1.** For each election system  $\mathcal{E}$  whose winner problem is in P, SR-online- $\mathcal{E}$ -UCM is in  $\Sigma_2^P$ .

---

<sup>1</sup>Sets  $S_1$  and  $S_2$  are said to be *P-separable* (see (Grollmann and Selman 1988)) if there exists a polynomial-time computable set  $T$  such that  $S_1 \subseteq T \subseteq \overline{S_2}$ . (One cannot in our main text change “are P-separable” into “are disjoint,” as then the reasoning used would become invalid; for example, SAT and  $\overline{\text{SAT}}$  are disjoint, and are respectively NP- and coNP-complete, but their union is  $\Sigma^*$  and so unless  $P = NP$  will not be  $P^{NP[1]}$ -complete.)

2. There exists an election system  $\mathcal{E}$ , whose winner problem is in P, such that SR-online- $\mathcal{E}$ -UCM (indeed, even SR-online- $\mathcal{E}$ -UCM[1, 1]) is  $\Sigma_2^P$ -complete.

The proof of Theorem 13 is omitted due to space limitations.

## Conclusions and Open Questions

We introduced a novel framework for online manipulation in sequential voting, and showed that manipulation there can be tremendously complex even for systems with simple winner problems. We also showed that among the most important election systems, some have efficient online manipulation algorithms but others (unless P = NP) do not. It will be important to, complementing our work, conduct typical-case complexity studies. Also, we will extend the scope of our investigation by studying online control and online bribery in appropriate models.

## References

- [2010] Bachrach, Y.; Betzler, N.; and Faliszewski, P. 2010. Probabilistic possible winner determination. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 697–702. AAAI Press.
- [1991] Bartholdi, III, J., and Orlin, J. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8(4):341–354.
- [1989] Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.
- [2012] Baumeister, D., and Rothe, J. 2012. Taking the final step to a full dichotomy of the possible winner problem in pure scoring rules. *Information Processing Letters* 112(5):186–190.
- [2010] Betzler, N., and Dorn, B. 2010. Towards a dichotomy of finding possible winners in elections based on scoring rules. *Journal of Computer and System Sciences* 76(8):812–836.
- [2009] Betzler, N.; Hemmann, S.; and Niedermeier, R. 2009. A multivariate complexity analysis of determining possible winners given incomplete votes. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 53–58. AAAI Press.
- [2010] Betzler, N. 2010. On problem kernels for possible winner determination under the k-approval protocol. In *Proceedings of the 35th International Symposium on Mathematical Foundations of Computer Science*, 114–125. Springer-Verlag Lecture Notes in Computer Science #6281.
- [1998] Borodin, A., and El-Yaniv, R. 1998. *Online Computation and Competitive Analysis*. Cambridge University Press.
- [2007] Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3):Article 14.
- [2010] Conitzer, V. 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53(3):84–94.
- [2010] Desmedt, Y., and Elkind, E. 2010. Equilibria of plurality voting with abstentions. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, 347–356. ACM Press.
- [2001] Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, 613–622. ACM Press.
- [2010] Faliszewski, P., and Procaccia, A. 2010. AI’s war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.
- [2009] Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2009. A richer understanding of the complexity of election systems. In Ravi, S., and Shukla, S., eds., *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*. Springer. 375–406.
- [2010] Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2010. Using complexity to protect elections. *Communications of the ACM* 53(11):74–82.
- [2008] Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2008. Copeland voting: Ties matter. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, 983–990. International Foundation for Autonomous Agents and Multiagent Systems.
- [1999] Ghosh, S.; Mundhe, M.; Hernandez, K.; and Sen, S. 1999. Voting for movies: The anatomy of recommender systems. In *Proceedings of the 3rd Annual Conference on Autonomous Agents*, 434–435. ACM Press.
- [1973] Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41(4):587–601.
- [1988] Grollmann, J., and Selman, A. 1988. Complexity measures for public-key cryptosystems. *SIAM Journal on Computing* 17(2):309–335.
- [2007] Hemaspaandra, E., and Hemaspaandra, L. 2007. Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1):73–83.
- [2005] Konczak, K., and Lang, J. 2005. Voting procedures with incomplete preferences. In *Proceedings of the Multidisciplinary IJCAI-05 Workshop on Advances in Preference Handling*, 124–129.
- [1988] Krentel, M. 1988. The complexity of optimization problems. *Journal of Computer and System Sciences* 36(3):490–509.
- [2010] Maudet, N.; Lang, J.; Chevaleyre, Y.; and Monnot, J. 2010. Possible winners when new candidates are added: The case of scoring rules. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 762–767. AAAI Press.
- [1972] Meyer, A., and Stockmeyer, L. 1972. The equivalence problem for regular expressions with squaring requires exponential space. In *Proceedings of the 13th IEEE Symposium on Switching and Automata Theory*, 125–129. IEEE Press.
- [2011] Parkes, D., and Procaccia, A. 2011. Dynamic social choice: Foundations and algorithms. Working paper.
- [1975] Satterthwaite, M. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.
- [1993] Sloth, B. 1993. The theory of voting and equilibria in non-cooperative games. *Games and Economic Behavior* 5(1):152–169.
- [1976] Stockmeyer, L. 1976. The polynomial-time hierarchy. *Theoretical Computer Science* 3(1):1–22.
- [2004] Tennenholtz, M. 2004. Transitive voting. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, 230–231. ACM Press.
- [2008] Xia, L., and Conitzer, V. 2008. Determining possible and necessary winners under common voting rules given partial orders. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, 196–201. AAAI Press.
- [2010a] Xia, L., and Conitzer, V. 2010a. Stackelberg voting games: Computational aspects and paradoxes. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 697–702. AAAI Press.

- [2010b] Xia, L., and Conitzer, V. 2010b. Strategy-proof voting rules over multi-issue domains with restricted preferences. In *Proceedings of the 6th International Workshop On Internet And Network Economics*, 402–414. Springer-Verlag Lecture Notes in Computer Science #6484.
- [2010] Xia, L.; Conitzer, V.; and Lang, J. 2010. Aggregating preferences in multi-issue domains by using maximum likelihood estimators. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 399–406. International Foundation for Autonomous Agents and Multiagent Systems.
- [2011] Xia, L.; Conitzer, V.; and Lang, J. 2011. Strategic sequential voting in multi-issue domains and multiple-election paradoxes. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, 179–188. ACM Press.
- [2011] Xia, L.; Lang, J.; and Conitzer, V. 2011. Hypercubewise preference aggregation in multi-issue domains. In *Proceedings of the 22st International Joint Conference on Artificial Intelligence*, 158–163. AAAI Press.